

Multiparty Symmetric Sum Types

Lasse Nielsen

DIKU, University of Copenhagen

Nobuko Yoshida

Imperial College London

Kohei Honda

Queen Mary, University of London

Abstract This paper introduces a new theory of multiparty session types based on symmetric sum types, by which we can type non-deterministic orchestration choice behaviours. While the original branching type in session types can represent a choice made by a single participant and accepted by others determining how the session proceeds, the symmetric sum type represents a choice made by agreement among all the participants of a session. Such behaviour can be found in many practical systems, including collaborative workflow in healthcare systems for clinical practice guidelines (CPGs). Processes using the symmetric sums can be embedded into the original branching types using conductor processes. We show that this type-driven embedding preserves typability, satisfies semantic soundness and completeness, and meets the encodability criteria [18, 9] adapted to the typed setting. The theory leads to an efficient implementation of a prototypical tool for CPGs which automatically translates the original CPG specifications from a representation called the Process Matrix to symmetric sum types, type checks programs and executes them.

1 Introduction

Clinical Practice Guidelines (CPGs) [21] are detailed descriptions of medical treatment procedures, practised globally with local variations, in order to treat specific medical disorders. CPGs are an example of social interactions, which include workflow models and various cooperation models: its richness stems from the diverse collaborative patterns human organisations can exhibit. One such pattern, which plays a prominent role in CPGs, is *symmetric synchronisation* where all the participants are equal in the decision-making, i.e. the participants collectively decide on one of the possible choices.

Motivated from practice, this paper aims to distill the essence of this symmetric synchronisation as an interaction primitive, position it as part of the type theory for the asynchronous π -calculus with multiparty sessions, and explore its properties to model workflow frameworks, enjoying the richness of multiparty session types to express how data is exchanged. Our starting point is a widely known semi-formal modelling framework for CPGs and other workflows called Process Matrix [14], which provides a concise and general description of symmetric synchronisation patterns as found in CPGs.

The new synchronisation primitive is generally useful, also for other calculi and applications. We add the symmetric synchronisation primitive to the asynchronous π -calculus and study it in a typed setting because it allows us to model CPGs as types, and enables correctness and erasure properties.

We explain the key ideas of Process Matrix and CPGs using an example from a CPG with three participants: A doctor (D), a nurse (N) and a patient (P). The doctor and the nurse need to register and inspect the patient, thus they must obtain the patient data (Data), schedule an appointment (Schedule) and inspect the

Figure 1.1 Cases in the healthcare cooperation example

	Data	Schedule	Inspect		
Case DD	D	D	D	D : Doctor N : Nurse P : Patient	Data: Obtain patient data Schedule: Schedule inspection Inspect: Perform inspection
Case ND	N	D	D		
Case DN	D	N	D		
Case NN	N	N	D		

patient (Inspect). The actions can be divided between the doctor and the nurse in four different ways, since they both can collect the data and schedule the appointment but only the doctor may inspect the patient. The four cases are illustrated in the table in Fig. 1.1. For example in Case ND, the nurse obtains the patient data and the doctor schedules and performs the inspection. In this way, the doctor and the nurse need to perform a different combination of actions depending on which case is chosen, thus they need to commit to the same choice, in order for the cooperation to work. This choice cannot be implemented directly using the asymmetric choice (as found in branching/selection primitives in the foregoing session types [20, 11]), since the decision would be done by a single participant and not by common agreement.

Our aim is to obtain a general modelling framework which can uniformly capture both symmetric synchronisations and existing session-based communication patterns. Such a framework will give a basis for the implementation of a tool for CPGs where one can describe, validate and execute specifications backed up by static validation coming from the theory. For this purpose we incorporate the synchronisation primitive in the type theory for multiparty sessions from [4, 12], so different groups of principals freely can mix standard asymmetric communications and symmetric synchronisations. The resulting sessions are abstracted as types, enabling type-based validation which ensures type and communication safety.

We offer the first prototype implementation of the π -calculus with multiparty sessions, with a type-checker using multiparty session types with full projections. Our implementation includes the symmetric synchronisation primitive and verification using symmetric sum types. This allows us to implement, verify and execute the examples used to explain and motivate the extension.

The use of types is not only essential for modelling CPGs and validating processes, but also enables an organised analysis of the synchronisation primitive. Using a type-directed translation, we show that the primitive can be embedded into the asymmetric branching in the original multiparty sessions [4, 12]. The translation generates auxiliary processes from the types, and combines them with an encoding of the sum into asymmetric branch types, respecting global interaction patterns and preserving semantics, by exploiting the type structure. The auxiliary process generated from a type conducts the synchronisations of a session by receiving accepted cases from participants and sending the chosen case back. To prove its correctness, we use a new technique based on derivations of the multiparty session typing. The resulting translation introduces exponentially more branching cases (e.g. 64 for the running example), demonstrating the practical usefulness of the symmetric sum for compact description as well as offering a formally founded distributed implementation strategy of the primitive.

Next we present the calculus for multiparty symmetric synchronisation (Section 2) and study its type theory (Section 3). We then define a type-directed encoding (Section 4) of the symmetric sum into the asynchronous multiparty session; and investigate its encodability criteria by adapting the framework from [18, 9] to the typed setting. Finally we present an application of the theory to the formal CPG verification (Section 5), with a prototype implementation available from [1]. The technical contributions include *subject reduction* (Theorem 3.2) and *type/semantic correctness* of the encoding (Theorems 4.1, 4.2 and 4.4). The implementation demonstrates the correctness, feasible implementability and significance of the new primitive. In particular, an automatic mapping from Process Matrix to global types (Section 5) shows the expressiveness of multiparty session types. Appendix in the full version [17] includes the omitted definitions, examples and proofs, though the paper can be read independently.

Figure 2.1 The process language

$P ::= \text{sync}_{\tilde{s},n}\{l : P_l\}_{l \in L}$	synchronisation	$s \triangleleft l; P$	label selection
$\mid \text{rand}\{P_i\}_{i \in I}$	random choice	$s \triangleright \{l : P_l\}_{l \in L}$	label branching
$\mid \overline{a}[2..n](\tilde{s}).P$	session request	$\text{if } e \text{ then } P \text{ else } Q$	conditional
$\mid a[p](\tilde{s}).P$	session accept	$P Q$	parallel
$\mid s!(\tilde{e});P$	value sending	0	inaction
$\mid s?(\tilde{x});P$	value reception	$(\nu n)P$	restriction
$\mid s!\langle\tilde{s}\rangle;P$	delegation	$\text{def } D \text{ in } P$	recursion
$\mid s?(\langle\tilde{s}\rangle);P$	reception	$X\langle\tilde{e}\tilde{s}\rangle$	process call
		$s : \tilde{h}$	message queue
$D ::= \{X_i(\tilde{x}_i\tilde{s}_i) = P_i\}_{i \in I}$	declarations	$v ::= a \mid \text{true} \mid \text{false}$	values
$e ::= v \mid x \mid e \text{ and } e' \mid \text{not } e \mid \text{rand}\{v_i\}_{i \in I} \mid \dots$	expressions	$h ::= l \mid \tilde{v} \mid \tilde{s}$	messages

2 Processes with Synchronisation

This section introduces the syntax (Fig. 2.1) of the asynchronous multiparty session π -calculus [12] with the new sync primitive, and the judgement $P \rightarrow P'$ (Fig. 2.2, where $e \downarrow v$ denotes the evaluation of the expression e to the value v) describing the small-step semantics for processes. The syntax defines the values: $\{v, w, \dots\}$, expressions: $\{e, e', \dots\}$ and processes: $\{P, Q, \dots\}$ from the sets of channel names: $\{a, b, \dots\}$, value variables: $\{x, y, \dots\}$, session channels: $\{s, t, \dots\}$, labels: $\{l, m, \dots\}$ and process variables: $\{X, Y, \dots\}$.

Session request, $\overline{a}[2..n](\tilde{s}).P$ initiates a session with channels \tilde{s} (where \tilde{s} denotes a vector $s_1 \dots s_n$) over the public channel a with the other $n - 1$ participants of shape $a[p](\tilde{s}).Q_p$ for p from 2 to n ([LINK] in Fig. 2.2). Asynchronous communication in an established session is performed by sending and receiving values ([SEND, RECV]), transferring a session using session delegation and reception ([DELEG, SREC]), and label selection and branching ([LABEL, BRANCH]), where the branching process offers a number of labels and the selecting process chooses one of them.

The new $\text{sync}_{\tilde{s},n}\{l : P_l\}_{l \in L}$ constructor is interpreted as the process participating in a plenum decision between all the n processes in the session \tilde{s} reaching a common decision h from L . Afterwards the process proceeds as described in P_h . In [SYNC] in Fig. 2.2, h in the premise denotes the common label. We also add the $\text{rand}\{P_i\}_{i \in I}$ constructor which randomly selects one of its branches ([RAND]). This primitive can be expressed using if and a random expression (hence it does not add expressiveness from [12]), but simplifies the erasure mapping in Section 4.

Figure 2.2 The reduction rules

[LINK]		[SEND]	
$\overline{a}[2..n](\tilde{s}).P_1 a[2](\tilde{s}).P_2 \dots a[n](\tilde{s}).P_n \rightarrow (\nu \tilde{s})(P_1 P_2 \dots P_n s_1 : \emptyset \dots s_m : \emptyset)$		$\tilde{e} \downarrow \tilde{v}$ $s!(\tilde{e}); P s : \tilde{h} \rightarrow P s : \tilde{h} \cdot \tilde{v}$	
[RECV]	[LABEL]	[BRANCH]	
$s?(\tilde{x}); P s : \tilde{v} \cdot \tilde{h} \rightarrow P[\tilde{v}/\tilde{x}] s : \tilde{h}$	$s \triangleleft l; P s : \tilde{h} \rightarrow P s : \tilde{h} \cdot l$	$s \triangleright \{l_i : P_i\}_{i \in I} s : l_j \cdot \tilde{h} \rightarrow P_j s : \tilde{h}$	$j \in I$
[DELEG]	[SREC]	[IFT]	
$s!\langle\tilde{t}\rangle; P s : \tilde{h} \rightarrow P s : \tilde{h} \cdot \tilde{t}$	$s?(\langle\tilde{t}\rangle); P s : \tilde{t} \cdot \tilde{h} \rightarrow P s : \tilde{h}$	$e \downarrow \text{true}$ $\text{if } e \text{ then } P \text{ else } Q \rightarrow P$	
[IFF]	[DEF]	[SCOP]	
$e \downarrow \text{false}$ $\text{if } e \text{ then } P \text{ else } Q \rightarrow Q$	$\tilde{e} \downarrow \tilde{v} \quad X\langle\tilde{x}\tilde{s}\rangle = P \in D$ $\text{def } D \text{ in } X\langle\tilde{e}\tilde{s}\rangle Q \rightarrow \text{def } D \text{ in } P[\tilde{v}/\tilde{x}] Q$	$P \rightarrow P'$ $(\nu n)P \rightarrow (\nu n)P'$	
[PAR]	[DEFIN]	[STR]	
$P \rightarrow P'$ $P Q \rightarrow P' Q$	$P \rightarrow P'$ $\text{def } D \text{ in } P \rightarrow \text{def } D \text{ in } P'$	$P \equiv P' \quad P \rightarrow Q \quad Q \equiv Q'$ $P' \rightarrow Q'$	
[RAND]	[SYNC]		
$j \in I$ $\text{rand}\{P_i\}_{i \in I} \rightarrow P_j$	$h \in \bigcap_{i=1}^n L_i$ $\text{sync}_{\tilde{s},n}\{l : P_{l_i}\}_{i \in L_1} \mid \dots \mid \text{sync}_{\tilde{s},n}\{l : P_{l_n}\}_{i \in L_n} \rightarrow P_{1h} \mid \dots \mid P_{nh}$		

Figure 2.3 Healthcare Example without sync

```

 $P_D$  = // Doctor
a[2](d, s, r, cp, cn).
if rand{true, false}
then
  cp<CaseD; cn<CaseD; d?(data);
  if rand{true, false}
  then
    cp<CaseDD; cn<CaseDD; s!(eSchedule); r!(eResult);end
  else
    cp<CaseDN; cn<CaseDN; r!(eResult);end
  else
    cp<CaseN; cn<CaseN;
    if rand{true, false}
    then
      cp<CaseND; cn<CaseND; s!(eSchedule); r!(eResult);end
    else
      cp<CaseNN; cn<CaseNN; r!(eResult);end

```

```

 $P_P$  = // Patient
 $\bar{a}[2..3](d, s, r, cp, cn). \text{pd}>
\{ \text{CaseD: } d!(eData); \text{cp}>
\{ \text{CaseDD: } s?(schedule); r?(result); \text{end},
\text{CaseDN: } s?(schedule); r?(result); \text{end} \},
\text{CaseN: } d!(eData); \text{pd}>
\{ \text{CaseND: } s?(schedule); r?(result); \text{end},
\text{CaseNN: } s?(schedule); r?(result); \text{end} \}
\}$ 
```

```

 $P_N$  = // Nurse
a[3](d, s, r, cp, cn). cn>
\{ \text{CaseD: } cn>
\{ \text{CaseDD: } \text{end},
\text{CaseDN: } s!(eSchedule); \text{end} \},
\text{CaseN: } d?(data); cn>
\{ \text{CaseND: } \text{end},
\text{CaseNN: } s!(eSchedule); \text{end} \} \}

```

In [SYNC], the processes cannot perform the synchronisation if they do not share some common label, in which case the processes will be stuck. We also need to know how many participants are in the session in order to know when the synchronisation can step; otherwise the processes will be stuck. The typing system introduced in the next section ensures that sync satisfies these two conditions.

Healthcare Cooperation (1): Processes We motivate the symmetric synchronisation using the example from the introduction. We first explain the problem when representing this interaction without sync. As explained in the introduction, there is no rigorous way to decide which of the four cases will occur, as well as who will be the principal decision maker: we could let the doctor non-deterministically decide between the cases, and then we obtain the processes in Fig. 2.3, if we are to use the processes from [12]: similarly we could let the nurse or even the patient decide. None of these representations captures the cooperation where the doctor, the nurse and the patient should reach a common decision, because *it is impossible to know who takes the initiative*. Another problem is that we need to specify the choices in P_D , which is best captured by non-deterministic expressions like rand.

Fig. 2.4 describes the same example using sync where the intended cooperation is directly modelled. The case is logically decided by two choices: first it is decided who receives the patient data, and then it is decided who schedules the inspection. Since these decisions are not necessarily made at the same time, the processes select the case using two sequential synchronisations.

Figure 2.4 Healthcare Example using sync

```

 $P_P$  = // Patient
 $\bar{a}[2..3](d, s, r). \text{sync}((d, s, r), 3)$ 
\{ \text{CaseD: } d!(eData); \text{sync}((d, s, r), 3)
\{ \text{CaseDD: } s?(schedule); r?(result); \text{end}, \text{CaseDN: } s?(schedule); r?(result); \text{end} \},
\text{CaseN: } d!(eData); \text{sync}((d, s, r), 3)
\{ \text{CaseND: } s?(schedule); r?(result); \text{end}, \text{CaseNN: } s?(schedule); r?(result); \text{end} \} \}

```

```

 $P_D$  = // Doctor
a[2](d, s, r). sync((d, s, r), 3)
\{ \text{CaseD: } d?(data); \text{sync}((d, s, r), 3) \{ \text{CaseDD: } s!(eSchedule); r!(eResult); \text{end}, \text{CaseDN: } r!(eResult); \text{end} \},
\text{CaseN: } \text{sync}((d, s, r), 3) \{ \text{CaseND: } s!(eSchedule); r!(eResult); \text{end}, \text{CaseNN: } r!(eResult); \text{end} \} \}

```

```

 $P_N$  = // Nurse
a[3](d, s, r). sync((d, s, r), 3)
\{ \text{CaseD: } \text{sync}((d, s, r), 3) \{ \text{CaseDD: } \text{end}, \text{CaseDN: } s!(eSchedule); \text{end} \},
\text{CaseN: } d?(data); \text{sync}((d, s, r), 3) \{ \text{CaseND: } \text{end}, \text{CaseNN: } s!(eSchedule); \text{end} \} \}

```

Figure 3.1 The Domains used for Global and Local types

(Global Types)	(Local Types)	(Message Types)
$G ::= p \rightarrow p' : k\langle U \rangle . G'$	$T ::= k!\langle U \rangle ; T$	$U ::= \tilde{S} \mid T@(\mathbf{p}, m, n)$
$\mid p \rightarrow p' : k\{l_i : G_i\}_{i \in I}$	$\mid k?\langle U \rangle ; T$	(Simple Types)
$\mid \mu t . G \mid t \mid \text{end}$	$\mid k \oplus \{l : T_l\}_{l \in L}$	$S ::= \text{bool} \mid \text{int} \mid \dots \mid \langle G \rangle$
$\mid \{l : G_l\}_{l \in L; M} (M \neq \emptyset)$	$\mid k \& \{l : T_l\}_{l \in L}$	(Environments)
	$\mid \mu t . T \mid t \mid \text{end}$	$\Gamma ::= \emptyset \mid \Gamma, u : \langle G \rangle \mid \Gamma, X : \tilde{S}\tilde{T}$
	$\mid \{l : T_l\}_{l \in L; M} (M \neq \emptyset)$	$\Delta ::= \emptyset \mid \Delta, \tilde{s} : T@(\mathbf{p}, n)$

3 Symmetric Sum Types

We start by defining the global types G in Fig. 3.1, which specifies global session protocols between the participants. Except for the symmetric sum type, the syntax is from [12]. The type $p \rightarrow p' : k\langle U \rangle . G'$ expresses that participant p sends a message of type U along channel k to p' and then interactions described in G' take place. The type $p \rightarrow p' : k\{l_i : G_i\}_{i \in I}$ expresses that p sends one of the labels l_i to p' . If l_j is sent, interactions described in G_j take place. Type $\mu t . G$ is a recursive type, assuming type variables (t, t', \dots) are guarded in the standard way. We assume that G in the grammar of sorts is closed, i.e., without free type variables. Type end represents the session termination.

The sum type $\{l : G_l\}_{l \in L; M}$ represents a synchronisation where the labels are taken from the set L and the non-empty set M . The labels in L are optional, but the labels in M are mandatory and must be accepted by all the participants. The mandatory labels will be underlined to distinguish them from the optional labels (e.g. $\{l : G_l\}_{l \in \{l_1\}; \{l_2\}} = \{l_1 : G_{l_1}, \underline{l_2} : G_{l_2}\}$).

The local types T are defined in Fig. 3.1. They describe the communication performed by a single process. Therefore the “from process to process on channel” syntax is simply changed to sending or receiving on a channel. Thus the sending type is $k!\langle U \rangle ; T$ and represents sending a message of type U on channel k , followed by the communication described by T . The type of receiving is $k?\langle U \rangle ; T$, the type of selecting is $k \oplus \{l : T_l\}_{l \in L}$ and the type of branching is $k \& \{l : T_l\}_{l \in L}$. The difference from [12] is that the symmetric sum type constructor $\{l : T_l\}_{l \in L; M}$ is added where L, M satisfies the conditions similar to those of a global sum type.

The message type $T@(\mathbf{p}, m, n)$ is used for delegation. It describes an open session, and includes information about the participant number \mathbf{p} , the number of session channels m , and the number of participants n in the session together with a local type T describing the remaining communication.

Finally we define the global environment Γ containing the global types for shared channels u , and process variables X , and the local type environment Δ containing the remaining session communication in Fig. 3.1, where $\tilde{s} : T@(\mathbf{p}, n)$ means \tilde{s} is an open session with n participants, where T describes the remaining communication for participant \mathbf{p} .

The *projection* $G|_{\mathbf{p}}$ of a global type G for a participant \mathbf{p} generates the local type for the participant in an intuitive way, for example $(p_0 \rightarrow p_1 : k\langle U \rangle . G')|_{\mathbf{p}}$ becomes $k!\langle U \rangle ; (G'|_{\mathbf{p}})$ if $\mathbf{p} = p_0$ and $\mathbf{p} \neq p_1$. The differences from the definition in [12] is that we have added a case for the symmetric sum type, $(\{l : G_l\}_{l \in L; M})|_{\mathbf{p}} = \{l : (G_l|_{\mathbf{p}})\}_{l \in L; M}$.

A global type G is coherent [12] if and only if the projection $G|_{\mathbf{p}}$ is defined for all participants, and G does not allow racing conditions (linearity). We only consider coherent global types.

Judgement The typing judgement extends the one from [12] with symmetric sum types. The judgement $\Gamma \vdash P \triangleright \Delta$ states that the process P in the environment Γ performs exactly the session communication described in Δ .

Figure 3.2 Selected typing rules

$\frac{[RAND] \quad \forall i \in I. \Gamma \vdash P_i \triangleright \Delta}{\Gamma \vdash \text{rand}\{P_i\}_{i \in I} \triangleright \Delta}$	$\frac{[SYNC] \quad \forall l \in L'' : \Gamma \vdash P_l \triangleright \Delta, \tilde{s} : T_l @ (p, n) \quad L'' \subseteq L \cup L' \quad L' \subseteq L''}{\Gamma \vdash \text{sync}_{\tilde{s}, n} \{l : P_l\}_{l \in L''} \triangleright \Delta, \tilde{s} : \{l : T_l\}_{l \in L; L'} @ (p, n)}$
$\frac{[MCAST] \quad \Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G 1) @ (1, n) \quad \tilde{s} = \max(\text{sid}(G)) \quad n = \max(\text{pid}(G))}{\Gamma \vdash \bar{a}[2..n](\tilde{s}).P \triangleright \Delta}$	
$\frac{[MACC] \quad \Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G p) @ (p, n) \quad \tilde{s} = \max(\text{sid}(G)) \quad n = \max(\text{pid}(G))}{\Gamma \vdash a[p](\tilde{s}).P \triangleright \Delta}$	
$\frac{[SEND] \quad \forall j. \Gamma \vdash e_j : S_j \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : T @ (p, n)}{\Gamma \vdash s_k!(\tilde{e}); P \triangleright \Delta, \tilde{s} : k!\langle \tilde{S} \rangle; T @ (p, n)}$	$\frac{[RCV] \quad \Gamma, \tilde{x} : \tilde{S} \vdash P \triangleright \Delta, \tilde{s} : T @ (p, n)}{\Gamma \vdash s_k?(\tilde{x}); P \triangleright \Delta, \tilde{s} : k?\langle \tilde{S} \rangle; T @ (p, n)}$
$\frac{[SEL] \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : T @ (p, n) \quad h \in L}{\Gamma \vdash s_k \triangleleft h; P \triangleright \Delta, \tilde{s} : k \oplus \{l : T_l\}_{l \in L} @ (p, n)}$	$\frac{[BRANCH] \quad \forall l \in L : \Gamma \vdash P_l \triangleright \Delta, \tilde{s} : T_l @ (p, n)}{\Gamma \vdash s_k \triangleright \{l : P_l\}_{l \in L} \triangleright \Delta, \tilde{s} : k? \{l : T_l\}_{l \in L} @ (p, n)}$
$\frac{[CONC] \quad \Gamma \vdash P \triangleright \Delta \quad \Gamma \vdash Q \triangleright \Delta'}{\Gamma \vdash P Q \triangleright \Delta \circ \Delta'} \quad (\text{dom}(\Delta) \cap \text{dom}(\Delta') = \emptyset)$	

The main rules are included in Fig. 3.2. The local types now carry information about the number of participants n and channels m . The number of participants and channels is determined at the session initialisation in the rules [MCAST] and [MACC], where $\text{sid}(G)$ denotes channels that appear in G and $\text{pid}(G)$ denotes the participants that appear in G . The rule [SYNC] checks that the synchronisation uses the correct number of participants, the accepted branches includes the mandatory ones and does not exceed the optional ones, and checks that each accepted branch is typed with the correct communication. The typing rule [RAND] checks that each choice in a rand process has the same session environment.

Since the process is reduced by each rule-application, the typability question $\Gamma \vdash P \triangleright \Delta$ is decidable.

Healthcare Cooperation (2): Types We explain how the types can describe and verify the healthcare scenario in the Introduction. Recall the processes from Fig. 2.4. To type $P_P | P_D | P_N$, we need a matching type-environment first. The processes use the public channel a to create a session, so the environment must be of the form $\Gamma = a : \langle G \rangle$ for some global type G .


We will start by finding the type describing the interactions in CaseND. First the participants select the choice CaseN and the patient sends the data to the nurse. Then the participants select the choice CaseND, the doctor sends the schedule to the patient, and finally the doctor sends the result to the patient.

When the patient has id 1, the doctor has id 2 and the nurse has id 3 the described communication for CaseND is described by the type

$\{\text{CaseN: } 1 \rightarrow 3 : 1 \langle \text{Sdata} \rangle . \quad \{\text{CaseND: } 2 \rightarrow 1 : 2 \langle \text{Sschedule} \rangle . \quad 2 \rightarrow 1 : 3 \langle \text{Sresult} \rangle . \text{end}\} \}$

Figure 3.3 Global Type G and Patient Projection for Healthcare Example

<pre> G = // Global type {CaseD: 1→2:1 ⟨Sdata⟩; {CaseDD: 2→1:2 ⟨Sschedule⟩; 2→1:3 ⟨Sresult⟩ ;end, CaseDN: 3→1:2 ⟨Sschedule⟩; 2→1:3 ⟨Sresult⟩ ;end }, CaseN: 1→3:1 ⟨Sdata⟩; {CaseND: 2→1:2 ⟨Sschedule⟩; 2→1:3 ⟨Sresult⟩ ;end, CaseNN: 3→1:2 ⟨Sschedule⟩; 2→1:3 ⟨Sresult⟩ ;end } } </pre>	<pre> G 1 = // Local type for Patient {CaseD: 1!⟨Sdata⟩; {CaseDD: 2?⟨Sschedule⟩; 3?⟨Sresult⟩ ;end, CaseDN: 2?⟨Sschedule⟩; 3?⟨Sresult⟩ ;end }, CaseN: 1!⟨Sdata⟩; {CaseND: 2?⟨Sschedule⟩; 3?⟨Sresult⟩ ;end, CaseNN: 2?⟨Sschedule⟩; 3?⟨Sresult⟩ ;end } } </pre>
---	--

Figure 4.1 Synchronisation message flows

Performing the same reasoning for CaseDD, CaseDN and CaseNN and adding their branches to the symmetric sums results in the global type G in Fig. 3.3. We select CaseND, CaseDN and CaseN as the mandatory labels. Since all participants must accept the mandatory choices, this means that it is always possible for the participants to agree on a choice in each of the synchronisations. We can then find the local type for the patient process as the patient's projection of G , given in Fig. 3.3. Using this type and the projections we can now typecheck the processes.

Proposition 3.1 $a : \langle G \rangle \vdash P_D \mid P_N \mid P_P \triangleright \emptyset$.

We end this section by proving *subject reduction*, from which we can derive soundness, communication safety and progress [12, § 5] as corollaries. Below $\Delta \rightarrow^{0/1} \Delta'$ denotes zero or one step using the type reduction [12], which represents the communication between dual local types. For instance, a reduction between input and output types is defined as: $k! \langle U \rangle; T_1 @ (p, n), k? \langle U \rangle; T_2 @ (q, n) \rightarrow T_1 @ (p, n), T_2 @ (q, n)$. We extend it to the symmetric sum as: $\{\{l : T_p, \dots\} @ (p, n)\}_{p \in \{1..n\}} \rightarrow \{T_p @ (p, n)\}_{p \in \{1..n\}}$.

The formulation uses the extension of the typing to runtime processes ($\Gamma \vdash P \triangleright_{\tau} \Delta$), which corresponds to the presented typing on processes without open sessions, but also accept processes with open sessions. This is obtained by joining compatible session environments (Δ, Δ') using the $\Delta \circ \Delta'$ operation to a single environment expressing the communication in both Δ and Δ' . Then we have:

Theorem 3.2 (Subject Reduction)

If $\Gamma \vdash P \triangleright_{\delta} \Delta$, Δ coherent and $P \rightarrow P'$ then $\Gamma \vdash P' \triangleright_{\delta} \Delta'$ where $\Delta \rightarrow^{0/1} \Delta'$.

PROOF: By induction on the derivation of $P \rightarrow P'$.

4 From Symmetric Sum to Conducted Branching

This section studies an erasure of symmetric synchronisation, which translates away symmetric sums using existing session primitives, which we hereafter simply call *the erasure*. The erasure removes all occurrences of the sync constructor while preserving static and dynamic semantics, i.e. typability and reduction. It uses a conductor process for each session. The messages and protocol used to implement the synchronisation are illustrated in Fig. 4.1 where the numbers indicate the sequence of the messages. Fig. 4.1(a) shows the communication between the processes without using sync in Fig. 2.3. Fig. 4.1(b) shows the communication between the processes using sync in Fig. 2.4, where no messages are sent, because the synchronisation ensures the same branch is chosen. Fig. 4.1(c) shows the conduction messages in the processes where the synchronisation has been erased in Fig. 4.5. First the patient, the doctor and the nurse send the cases they can accept to the conductor, who chooses a common case and sends the selected case to the patient, the doctor and the nurse.

Figure 4.2 Erasure of Synchronisation from Typing-Derivation

$$\begin{aligned}
\mathcal{E} \left[\frac{[\text{MCAST}] \quad \Gamma \vdash a : \langle G \rangle \quad \mathcal{D}_1 :: \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G|1)@(1, n) \quad \begin{array}{l} |\tilde{s}| = \max(\text{sid}(G)) \\ n = \max(\text{pid}(G)) \end{array}}{\Gamma \vdash \bar{a}[2..n](\tilde{s}).P \triangleright \Delta} \right] &= \mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a} \mid \mathcal{E} \llbracket \mathcal{D}_1 \rrbracket \\
\mathcal{E} \left[\frac{[\text{MACC}] \quad \Gamma \vdash a : \langle G \rangle \quad \mathcal{D}_1 :: \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G|p)@(p, n) \quad \begin{array}{l} |\tilde{s}| = \max(\text{sid}(G)) \\ n = \max(\text{pid}(G)) \end{array}}{\Gamma \vdash a[p](\tilde{s}).P \triangleright \Delta} \right] &= a[p](\tilde{s}, \text{in}_{\tilde{s}1}, \text{out}_{\tilde{s}1}, \dots, \text{in}_{\tilde{s}n}, \text{out}_{\tilde{s}n}). \mathcal{E} \llbracket \mathcal{D}_1 \rrbracket \\
\mathcal{E} \left[\frac{[\text{SYNC}] \quad \forall l \in L' : \mathcal{D}_l :: \Gamma \vdash P_l \triangleright \Delta, \tilde{s} : T_l @ (p, n) \quad L' \subseteq L \cup M \quad M \subseteq L'}{\Gamma \vdash \text{sync}_{\tilde{s}, n} \{l : P_l\}_{l \in L'} \triangleright \Delta, \tilde{s} : \{l : T_l\}_{l \in L, M, n} @ (p, n)} \right] &= \text{out}_{\tilde{s}p} \triangleleft \text{cases}_{L'} ; \text{in}_{\tilde{s}p} \triangleright \{l : \mathcal{E} \llbracket \mathcal{D}_l \rrbracket\}_{l \in L'} \\
\mathcal{E} \left[\frac{[\text{LABEL}] \quad \mathcal{D}_1 :: \Gamma \vdash P \triangleright \Delta, \tilde{s} : T_h @ (p, n) \quad h \in L}{\Gamma \vdash s_k \triangleleft h ; P \triangleright \Delta, \tilde{s} : k \oplus \{l : T_l\}_{l \in L} @ (p, n)} \right] &= s_k \triangleleft h ; \text{out}_{\tilde{s}p} \triangleleft h ; \mathcal{E} \llbracket \mathcal{D}_1 \rrbracket
\end{aligned}$$

The other cases are defined monomorphic

4.1 Erasure Definitions

Based on this idea, we translate the synchronisation and symmetric sum types into the original system [12], step by step as follows.

Step 1: Process Erasure Only well-typed processes are eligible for erasure, because conductor processes are generated from the global types. Therefore the erasure $\mathcal{E} \llbracket \cdot \rrbracket$ is defined on the type derivation in Fig. 4.2 and the result is the erased process. We use the notation $\mathcal{D} :: \Gamma \vdash P \triangleright \Delta$ to denote a derivation \mathcal{D} with the conclusion $\Gamma \vdash P \triangleright \Delta$.

The case for session request increments the number of participants by one, to make room for the conductor process, and adds two session channels per user ($\text{in}_{\tilde{s}, p}$ and $\text{out}_{\tilde{s}, p}$), for communicating with the conductor. The conductor process $\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a}$ (defined in Step 2) is inserted in parallel with the resulting session requesting process to ensure it is available.

The case for synchronisation sends the accepted labels to the conductor, waits to receive one of the accepted labels and proceeds with the selected branch.

Step 2: Conductor Generation The conductor process $\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a}$ was inserted in parallel with the session requests by the process erasure in Step 1. The main cases of the conductor generation $\mathcal{C} \llbracket \cdot \rrbracket$ are in Fig. 4.3. Notice that $\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a}$ is only a wrapper for $\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n}^*$ which prefixes the session acceptance on channel a . In $\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a}$, \tilde{s} is the original session channels, n is the number of original participants, G is the original session type, and a is the channel the session is created over.

The conductor process generated from a synchronisation receives the accepted labels from each participant, selects a common label using `rand` and sends the selected label back to each participant before conducting the chosen branch.

Step 3: Type Translations To prove that typability is preserved by the erasure, we define translations of global types, local types, message types, global type environments and local type environments to find the types for the result of the erasure. The main cases for global types are defined in Fig. 4.4. The translation $\llbracket G \rrbracket$ of global types is just a wrapper for $\llbracket G \rrbracket_{n, m}^*$ where n is the number of participants, and m is the number of session channels in the original type.

Figure 4.3 Conductor Process Generation from a Global Type

$$\begin{aligned}
\mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n, a} &= a[n+1](\tilde{s}, \text{in}_{\tilde{s}1}, \text{out}_{\tilde{s}1}, \dots, \text{in}_{\tilde{s}n}, \text{out}_{\tilde{s}n}). \mathcal{C} \llbracket G \rrbracket_{\tilde{s}, n}^* \\
\mathcal{C} \llbracket \{l : G_l\}_{l \in L; M} \rrbracket_{\tilde{s}, n}^* &= \text{out}_{\tilde{s}1} \triangleright \{ \text{cases}_{L_1 \cup M} : \dots : \text{out}_{\tilde{s}n} \triangleright \{ \text{cases}_{L_n \cup M} : \\
&\quad \text{rand} \{ \text{in}_{\tilde{s}1} \triangleleft l ; \dots ; \text{in}_{\tilde{s}n} \triangleleft l ; \mathcal{C} \llbracket G_l \rrbracket_{n, \tilde{s}}^* \}_{l \in \bigcap_{i=1}^n L_i \cup M} \}_{L_n \subseteq L \dots} \}_{L_1 \subseteq L}
\end{aligned}$$

Figure 4.4 Erasure Mapping for Global Types

$$\begin{aligned}
\llbracket G \rrbracket &= \llbracket G \rrbracket_{\max(\text{pid}(G)), \max(\text{sid}(G))}^* \\
\llbracket \{l : G_l\}_{l \in L; M} \rrbracket_{n, m}^* &= 1 \rightarrow n+1 : (m+2) \{ \text{cases}_{L_1 \cup M} : \\
&\quad 2 \rightarrow n+1 : (m+4) \{ \text{cases}_{L_2 \cup M} : \dots \\
&\quad n \rightarrow n+1 : (m+2 \cdot n) \{ \text{cases}_{L_n \cup M} : \\
&\quad n+1 \rightarrow 1 : (m+1) \{ l : n+1 \rightarrow 2 : (m+3) \{ l : \dots \\
&\quad n+1 \rightarrow n : (m+2 \cdot n - 1) \{ l : \llbracket G_l \rrbracket_{n, m}^* \dots \} \}_{l \in \bigcap_{i=0}^n L_i \cup M} \}_{L_n \subseteq L \dots} \}_{L_1 \subseteq L}
\end{aligned}$$

As previously suggested, the symmetric sum is translated to nested branching, where each participant sends the accepted labels to the conductor, receives the selected label and continues with the selected branch.

4.2 Correctness

We now prove the correctness of the erasure mapping. We start by proving that the typing is preserved, and the types of the result process is given by the defined type translations.

Theorem 4.1 (Type Preservation) *If $\mathcal{D} :: \Gamma \vdash P \triangleright \Delta$ then $\llbracket \Gamma \rrbracket \vdash \mathcal{E} \llbracket \mathcal{D} \rrbracket \triangleright \llbracket \Delta \rrbracket$*

PROOF: By induction on the type derivation \mathcal{D} . The proof uses a lemma stating that the generated conductor processes are well-typed.

Next we prove that process congruence ($P \equiv Q$) is preserved by the erasure.

Theorem 4.2 (Congruence Preservation)

If $\mathcal{D}_1 :: \Gamma \vdash P \triangleright_{\bar{\tau}} \Delta$ then for all Q we have that $P \equiv Q$ if and only if there is a derivation $\mathcal{D}_2 :: \Gamma \vdash Q \triangleright_{\bar{\tau}} \Delta$ such that $\mathcal{E} \llbracket \mathcal{D}_1 \rrbracket \equiv \mathcal{E} \llbracket \mathcal{D}_2 \rrbracket$.

Congruence preservation suggests the erasure preserves semantic properties. We start by stating the soundness theorem. To do this we define conductors for partially completed sessions: $PC(\Delta)$ as the set of possible partial conductor processes generated from Δ . By using the partial conductors from the session environment it is now possible to state the soundness theorem.

Theorem 4.3 (Soundness) *If $\mathcal{D} :: \Gamma \vdash P \triangleright_{\bar{\tau}} \Delta$, $P \rightarrow P'$, Δ coherent and $P_C \in PC(\Delta \circ \Delta'')$ for some Δ'' then there is a derivation $\mathcal{D}' :: \Gamma \vdash P' \triangleright_{\bar{\tau}} \Delta'$ and $P'_C \in PC(\Delta' \circ \Delta'')$ such that $\Delta \rightarrow^{0/1} \Delta'$ and $\mathcal{E} \llbracket \mathcal{D} \rrbracket | P_C \rightarrow^* \mathcal{E} \llbracket \mathcal{D}' \rrbracket | P'_C$.*

PROOF: By induction on the derivation of $P \rightarrow P'$.

We can extend the above theorem to multiple steps by induction on the number of steps. Also the found evaluation of $\mathcal{E} \llbracket \mathcal{D} \rrbracket \rightarrow^* \mathcal{E} \llbracket \mathcal{D}' \rrbracket$ performs exactly the same communication on all non-conductor channels as the original evaluation $P \rightarrow^* P'$.

We will now define *conduction steps*, since they play an important role in formulating the completeness theorem. This is because all steps performed by the result of the erasure can be mimicked by the original process up to *conduction steps*. A step from P_1 to P_2 is a conduction step, written $P_1 \rightarrow P_2$ if the step performs *label selection* or *label branching* on a conductor channel or unfolding of a *conductor process*; otherwise we write $P_1 \rightarrow P_2$. We observe all the extra steps introduced by the erasure are of the form \rightarrow , while the other steps are of the form \rightarrow . Therefore there is a one-to-one correspondence between the \rightarrow steps of the erased process, and the steps in the original process.

Theorem 4.4 (Semantic Completeness) *If $\mathcal{E} \llbracket \mathcal{D}_1 :: \Gamma \vdash P_1 \triangleright \emptyset \rrbracket \rightarrow^* Q'$ then there exists a derivation $\mathcal{D}_2 :: \Gamma \vdash P_2 \triangleright \emptyset$ and Q such that $P_1 \rightarrow^* P_2$ and $\mathcal{E} \llbracket \mathcal{D}_2 \rrbracket \rightarrow^* Q$ and $Q' \rightarrow^* Q$.*

Figure 4.5 Example Processes after Erasure

```

 $P'_C$  = // Conductor
a[4](d,s,r,in_p,out_p,
    in_d,out_d,in_n,out_n).
out_p ▷
{cases_DN: out_d ▷
 {cases_DN: out_n ▷
 {cases_DN:
   rand {
     in_p < CaseD;
     in_d < CaseD;
     in_n < CaseD;
     out_p ▷
     {cases_DN: out_d ▷
      {cases_DN: out_n ▷
       {cases_DN:
        rand
        {in_p < CaseDD;
         in_d < CaseDD;
         in_n < CaseDD;
         end,
         in_p < CaseDN;
         in_d < CaseDN;
         in_n < CaseDN;
         end },
         cases_D: ... },
         cases_D: ... },
         cases_D: ... } },
         cases_N: ... },
         cases_N: ... },
         cases_N: ... }

```

```

 $P'_P$  = // Patient
a[2..4](d,s,r,in_p,out_p,
    in_d,out_d,in_n,out_n).
out_p < cases_DN; in_p ▷
{CaseD: d!(eData); out_p < cases_DN; in_p ▷
 {CaseDD: s?(schedule); r?(result);0,
  CaseDN: s?(schedule); r?(result);0},
 CaseN: d!(eData); out_p < cases_DN; in_p ▷
 {CaseND: s?(schedule); r?(result);0,
  CaseNN: s?(schedule); r?(result);0} }

```

```

 $P'_D$  = // Doctor
a[2..4](d,s,r,in_p,out_p,
    in_d,out_d,in_n,out_n).
out_d < cases_DN; in_d ▷
{CaseD: d?(data); out_d < cases_DN; in_d ▷
 {CaseDD: s!(eSchedule); r!(eResult);0,
  CaseDN: r!(eResult);0},
 CaseN: out_d < cases_DN; in_d ▷
 {CaseND: s!(eSchedule); r?(eResult);0,
  CaseNN: r?(eResult);0} }

```

```

 $P'_N$  = // Nurse
a[2..4](d,s,r,in_p,out_p,
    in_d,out_d,in_n,out_n).
out_n < cases_DN; in_n ▷
{CaseD: out_n < cases_DN; in_n ▷
 {CaseDD: 0,
  CaseDN: s!(eSchedule);0},
 CaseN: d?(data); out_n < cases_DN; in_n ▷
 {CaseND: 0,
  CaseNN: s!(eSchedule);0} }

```

PROOF: By induction on the number of non-conduction steps in $\mathcal{E}[\llbracket \mathcal{D}_1 \rrbracket] \rightarrow^* Q'$, using confluence and single-step completeness results.

Healthcare Cooperation (3): Synchronisation Erasure The result of the erasure on the healthcare example from Section 3 is shown in Fig. 4.5. Since we have shown that the processes from the synchronisation example in Fig. 2.4 are well-typed in Proposition 3.1, we can apply Theorem 4.1 to provide $a : \langle \llbracket G \rrbracket \rangle \vdash P'_C \mid P'_P \mid P'_D \mid P'_N \triangleright \emptyset$.

As this example illustrates, the result of the erasure does not capture the nature of the situation in the same way, because it introduces a conductor process, which is not a natural part of the situation. It is not compact either, as the conductor process has 64 cases. Further we lose an accurate type abstraction of the dynamics of symmetric synchronisation, because it is not clear from the encoded type structure whether it is just a sequence of asymmetric branching actions or the (intended) atomic multiparty synchronisation, since some of the key operational structures of the encoding (e.g. random selection) is lost in the encoded type.

4.3 Encodability Criteria

The common properties of encodability from the known separation theorems (e.g. [18]) has been studied [9], revealing a number of desirable criteria. Our encoding is *type-based*, so we cannot apply this untyped framework directly. However if we simply change the formulation to use the *type-derivation* instead of the process syntax, our encoding *does* fulfil the criteria.

Before we can define and prove the criteria, we need to define the relations (\approx_1 and \approx_2) and properties (successful state) used to define the criteria. We select \approx_1 as the process equivalence (\equiv), and define $Q_1 \approx_2 Q_2$ if and only if $\exists Q. Q_1 \rightarrow^* Q \wedge Q_2 \rightarrow^* Q$.

Lemma 4.5 \asymp_2 is a weak barbed reduction congruence.

PROOF: Immediately \asymp_2 is symmetric and reflective by definition. By the confluence, we can also prove its transitivity.

To define a successful state, we introduce a new process constructor \surd , and extend the typing system to accept \surd , and extend the erasure to preserve \surd . A process P is accepting if $P \equiv \surd | P'$ for some P' .

We list the new formulation for all the criteria and state the theorem. For the motivation of each criterion, see [9]. Below, for the sake of readability, we omit Γ and Δ from the encoding.

Compositionality criterion For every k -ary typing rule R in the typing system of \mathcal{L}_1 and every subset of names N there exists a k -ary context $C_R^N(_, \dots, _)$ such that, for all $\mathcal{D}_1, \dots, \mathcal{D}_k$ with $\text{FN}(\mathcal{D}_1, \dots, \mathcal{D}_k) = N$, it holds that $\llbracket R(\mathcal{D}_1, \dots, \mathcal{D}_k) \rrbracket = C_R^N(\llbracket \mathcal{D}_1 \rrbracket, \dots, \llbracket \mathcal{D}_k \rrbracket)$. Note that the information given by derivation (typing) in $\mathcal{D}_1 :: P_1$ and $\mathcal{D}_2 :: P_2$ are essential.

Name Invariance criterion For every typing derivation $\mathcal{D} :: P$ (P has derivation \mathcal{D}) and name substitution σ , it holds that if σ is injective, then $\llbracket \mathcal{D}\sigma \rrbracket = \llbracket \mathcal{D} \rrbracket \sigma'$; for every $a \in \mathcal{N}$, otherwise $\llbracket \mathcal{D}\sigma \rrbracket \asymp_2 \llbracket \mathcal{D} \rrbracket \sigma'$ where σ' is such that $\varphi_{\llbracket \cdot \rrbracket}(\sigma(a)) = \sigma'(\varphi_{\llbracket \cdot \rrbracket}(a))$. Here $\varphi_{\llbracket \cdot \rrbracket}$ is called the renaming policy and captures how $\llbracket \cdot \rrbracket$ translates channel names.

Operational Correspondence criterion Let \rightarrow_i denote the reduction relation of the system i .

(1) **Completeness:** If $\mathcal{D}_1 :: P_1$ and $P_1 \rightarrow_1^* P_2$ then there exists a $\mathcal{D}_2 :: P_2$ such that $\llbracket \mathcal{D}_1 \rrbracket \rightarrow_2^* \asymp_2 \llbracket \mathcal{D}_2 \rrbracket$.

(2) **Soundness:** If $\llbracket \mathcal{D}_1 :: P_1 \rrbracket \rightarrow_2^* Q_1$ then there exists a $\mathcal{D}_2 :: P_2$ such that $P_1 \rightarrow_1^* P_2$ and $Q_1 \rightarrow_2^* \asymp_2 \llbracket \mathcal{D}_2 \rrbracket$.

Divergence Reflection criterion If $\llbracket \mathcal{D} :: P \rrbracket \rightarrow^\omega$ then $P \rightarrow^\omega$ where \rightarrow^ω means infinite reductions.

Success Sensitiveness criterion If $\mathcal{D} :: P$ then $P \Downarrow$ if and only if $\llbracket \mathcal{D} \rrbracket \Downarrow$ where $P \Downarrow$ means P can reach a successful state.

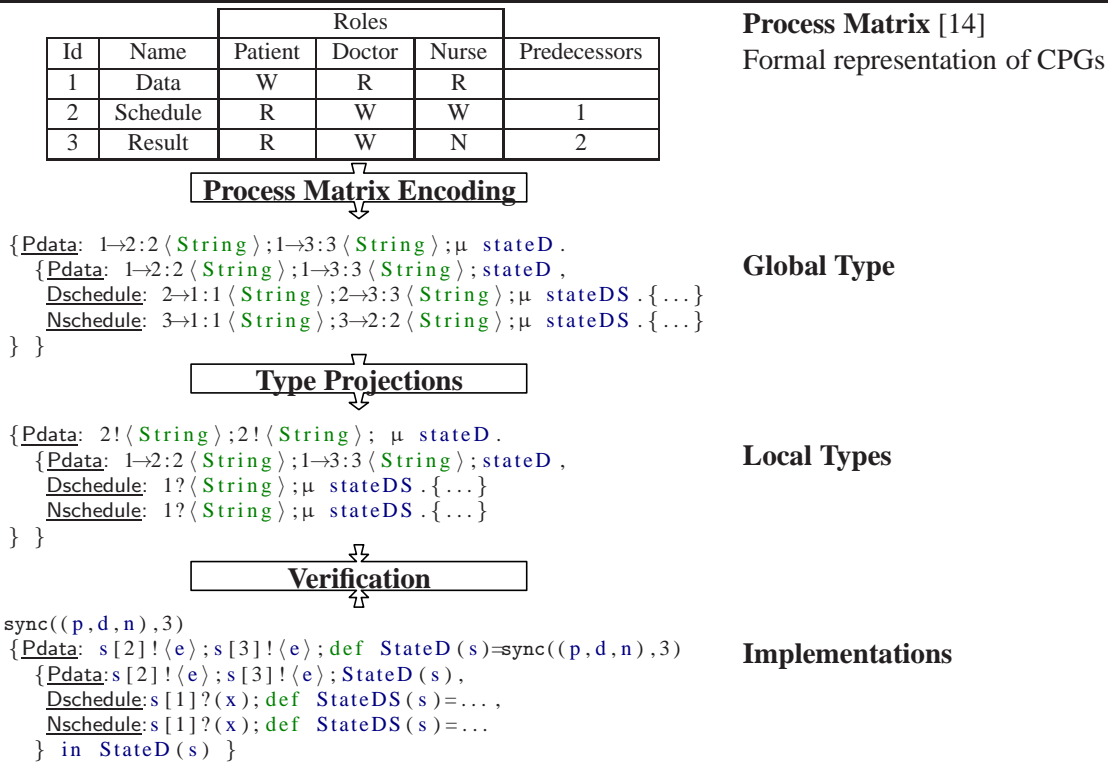
Using the above definition, we arrive at the following main theorem.

Theorem 4.6 The erasure mapping satisfies all the encodability criteria.

5 Verifying CPG Descriptions

This section describes how symmetric sum types can verify implementation conformance to a CPG [21] described using the Process Matrix. The verification is performed by three steps in Fig. 5.1, as illustrated below.

Process Matrix. The Process Matrix representation consists of a table with one row for each action. Each row has a number of columns: The **Id** and **Name** columns are used to identify the action, and the **Predecessors** column holds the *Ids* of the actions the action depends on. Before an action can be executed its predecessors must have been executed. If all the predecessors of an action have been executed we say that the action is *executable*. Finally there is one column for each participant (called roles), where the content is either R meaning the participant can read the action-data but not execute it, W meaning the participant can execute the action and read its data or N meaning the participant cannot execute the action or read its data (see [14] for a more adequate description). The Process Matrix in Fig. 5.1 describes the scenario from the introduction, except that the patient automatically gives the data to both the doctor and the nurse, and the user can perform the actions multiple times (by an implicit recursion), until all the actions are executed.

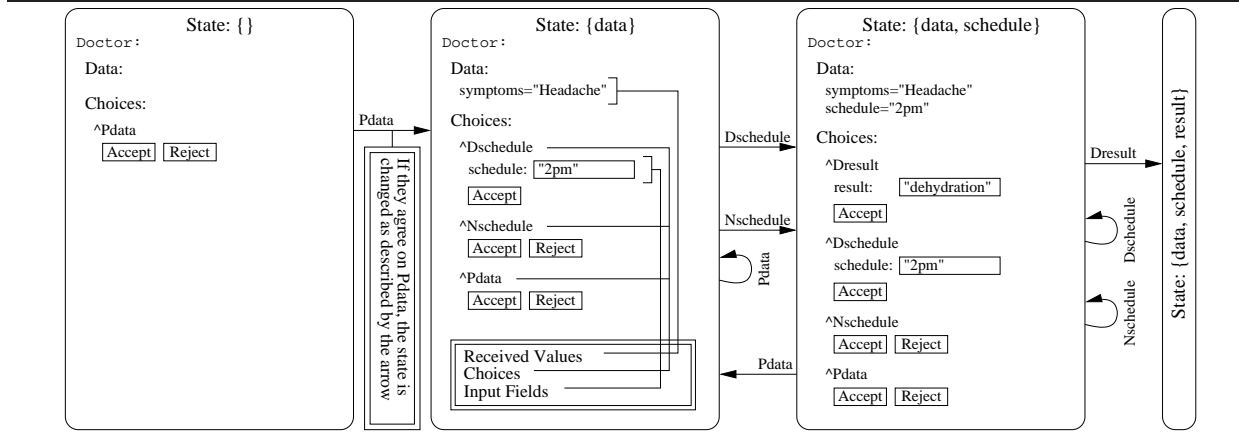
Figure 5.1 Steps in verifying a CPG description

Process Matrix Encoding Any CPG in a Process Matrix can be encoded as a global type automatically. We explain this encoding by translating the above Process Matrix example. In the resulting type, the state is described by the set of actions that have been executed, leading to a finite but exponential number of states. The representation of each state (except the completed state) is a symmetric sum with one branch for each role that can execute each executable action. The content of each branch consists of the executing participant sending the created data to all other participants with read or write access, followed by the state where the executed action is added, and depending actions have been removed.

Parts of the global type is included in Fig. 5.1. Notice that the resulting type uses recursion: this is to describe an implicit recursion in the Process Matrix where the state reached after an action does not have to be a new state, but can be the same as the state before the execution of the action, or even from previous steps. This is the case for the above example if the data is sent, the appointment is scheduled, and then the data is resent. The resulting state would then be the state where only the data action has been executed, which is the same as the second state. The described method can be extended to translate any Process Matrix into a global type.

The conversion of CPGs from the Process Matrix, to session type allows the data to be exchanged directly between the participants, while the current implementations rely on a centralised database for the exchange. This means the translation offers a distributed implementation of the Process Matrix, which has not been known before. A formally defined symmetric global synchronisation primitive, together with its type discipline and encodability, offers a firm basis for such implementations.

Projection and Verification When we have created the global type expressing the CPG, a process implementing one of the participants can be verified to conform with the workflow, by projecting the global type to the local type of that participant, and typechecking the process against the local type. Parts of the local type and the process for the Patient are described in Fig. 5.1.

Figure 5.2 States and screenshots for the doctor GUI

Generalisation We have now described how to use the multiparty session types extended with symmetric sum, to express CPGs formalised using the Process Matrix. We believe many other workflow frameworks (such as large parts of the BPMN) can be encoded as multiparty session types with symmetric sum, and this would allow the type-system to serve as a common representation, enabling interaction between different frameworks and implementing features (such as automatic user-interface generation) only for symmetric sum types, and apply it to all the encoded frameworks.

5.1 Implementation

We have created an `ascii` syntax for the asynchronous π -calculus with multiparty sessions and symmetric synchronisation called APIMS, and implemented a typechecker and an interpreter. This is to our knowledge the first prototype implementation of the π -calculus with multiparty sessions and multiparty session types. The implementation along with example programs can be found on the APIMS website [1].

The implementation extends the calculus with a `guisync` constructor to support user interaction via GUIs. The `guisync` is the result of extending the `sync` for user input. Each label has a set of typed arguments that must be given using the GUI before that choice is accepted, and the given arguments can be used by the process in that branch. This simple extension allows the processes to implement GUIs and the type system guarantees that the GUI for each participant will respect the protocol, hence the workflow. The mandatory labels ensure that the GUI must allow all the users (the people using the interface for each participant) to agree in each synchronisation, thus avoiding the GUIs causing a disagreement w.r.t. the theory of a symmetric synchronisation.

The GUI shows the received data, the choices offered by the process, input fields for the data needed for each choice, and buttons to accept/reject each choice. Fig. 5.2 shows three screen-shots, displaying the doctor's GUI for each state and how each choice affects the state. As soon as all the participants of a session accepts the same choice, the processes continue with the accepted branch. The GUI implementation for each participant can be created automatically from the Process Matrix.

The original implementation of the Process Matrix called *Online Consultant* by *Resultmaker* [14] is database based. This means that communication consists of the sender uploading information to the server, and all participants must query the server when using the information. Implementing the workflows using the π -calculus and session types not only gives the *Process Matrix* a formal semantics, but also allows an implementation where participants communicate their data as peer-to-peer. This offers more natural and robust realisation of the workflows, and relieves the system from the server bottleneck.

6 Related and Future Work

There are existing studies on self/broadcast synchronisations [10, 19]. The symmetric sum proposed in the present paper is different because it allows all the participants to influence the choice equally and, to formulate this notion adequately, demands a session-based operational framework. Another difference is the use of the type discipline to control this complex synchronisation framework, which is not found in the foregoing work. Note that the type discipline allows multiparty progress and communication-safety for participants, which is not generally ensured in existing untyped self/broadcast synchronisation primitives. Our primitive and its type-checker are applicable not only to Process Matrix, but also multiparty synchronisations in general with strong safety guarantees.

The symmetric synchronisation is similar to the consensus in Weak Byzantine Agreement (WBA) [7, 13, 2, 8] which is a formalisation of the database commit problem. The similarity is that a number of processes need to end up with a common choice. In contrast to symmetric sum, WBA only has two possible choices (0 and 1). Not all participant has to initially accept the final decision, but if all processes agree initially, the result should be the initial preference. WBA is studied in an untyped settings on unreliable networks, with faulty processes (with arbitrary behaviour).

The symmetric sum is also similar to the symmetric choice \square in CSP and the mixed choice in the π -calculus [18]. The main difference is these preceding primitives are restricted to two party synchronisations. Our result is consistent with the non-encodability of the mixed-choice π -calculus in the separated choice π -calculus [18]: our erasure is defined on *typing derivations*, and cannot be made homomorphic on *processes*. For example, take $P = (va)(P_1|P_2)$ where

$$P_1 = \bar{a}[2](s).\text{sync}\{\underline{l1} : P_{11}, \underline{l2} : P_{12}\} \text{ and } P_2 = a[2](s).\text{sync}\{\underline{l1} : P_{21}, \underline{l3} : P_{23}\}.$$

This process shows that the erasure cannot be interpreted as an encoding from processes $\llbracket \cdot \rrbracket$ where $\llbracket P_1|P_2 \rrbracket = \llbracket P_1 \rrbracket \llbracket P_2 \rrbracket$, because the result of $\llbracket P_1 \rrbracket$ depends on the context P_1 is in: the conductor inserted by the second step of the erasure depends on the type of a which depends on the other process. In the given context, the conductor must consider the labels $\underline{l1}, \underline{l2}$ and $\underline{l3}$, and this could not be generated from $\llbracket P_1 \rrbracket$ because P_1 does not contain any information about $\underline{l3}$. As noted above, the symmetric sum and synchronisation construct differs from the mixed choice and from the untyped asymmetric, directed sums whose encodability is studied in [16, 15], in that it is multi-party synchronisation for a fixed number of participants ensured by the underlying session type discipline.

Types for the multiparty interactions are studied in the *conversation calculus* [5] and *contracts* [6]. The former has choice behaviours where the channel-based communication is replaced by conversation environments allowing multiple participants, while the latter uses a process-based specification of protocols relying on internal and external choices, where conformance is formalised based on must preorder (so that we can ensure liveness). Our implementation crucially relies on the choreographic description based on global types: in particular, global types can offer a tractable, clear type-directed generation from the Process Matrices as described in Section 5.

As future work, we plan to extend our work with logical assertions based on [3] in order to describe and ensure the communicated data fulfil desired properties (for example, “the prescribed medicine doses are less than the lethal amount”). With the assertions, we can add arguments (state) to the recursive types, and conditions to the branches in a choice, so that it will lead to a more efficient generation from the Process Matrix.

Acknowledgements

The first author is supported by the *TrustCare* project, funded by the Danish Strategic Research Agency, Grant #2106-07-0019. The last two authors are partially supported by EPSRC EP/F003757, EP/F002114, EP/G015635 and EP/G015481.

References

- [1] Apims Project Page. Available at: <http://www.thelas.dk/index.php/apims>.
- [2] C. Attiya, D. Dolev, and J. Gil. Asynchronous byzantine consensus. In *PODC '84: Proceedings of the third annual ACM symposium on Principles of distributed computing*, pages 119–133, New York, NY, USA, 1984. ACM.
- [3] L. Bocchi, K. Honda, E. Tuosto, and N. Yoshida. A theory of design-by-contract for distributed multiparty interactions, 2009. To appear in CONCUR 2010.
- [4] E. Bonelli and A. B. Compagnoni. Multipoint session types for a distributed calculus. In *TGC*, volume 4912 of *LNCS*, pages 240–256. Springer, 2007.
- [5] L. Caires and H. T. Vieira. Conversation types. In *ESOP '09*, pages 285–300, Berlin, Heidelberg, 2009. Springer.
- [6] G. Castagna and L. Padovani. Contracts for mobile processes. In *CONCUR '09*, *LNCS*, pages 211–228, Berlin, Heidelberg, 2009. Springer.
- [7] M. J. Fischer. The consensus problem in unreliable distributed systems (a brief survey). In *Proceedings of the 1983 International FCT-Conference on Fundamentals of Computation Theory*, pages 127–140, London, UK, 1983. Springer.
- [8] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [9] D. Gorla. Towards a Unified Approach to Encodability and Separation Results for Process Calculi. In *CONCUR '08*, *LNCS*, pages 492–507, 2008.
- [10] T. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [11] K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, volume 1381 of *LNCS*, pages 22–138. Springer, 1998.
- [12] K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL'08*, pages 273–284. ACM, 2008.
- [13] L. Lamport. The weak byzantine generals problem. *J. ACM*, 30(3):668–676, 1983.
- [14] K. Lyng, T. Hildebrandt, and R. Mikkamala. From paper based clinical practice guidelines to declarative workflow management. In *LNBIP '09*, pages 336–347.
- [15] U. Nestmann. What is a “good” encoding of guarded choice? *Inf. Comput.*, 156(1-2):287–319, 2000.
- [16] U. Nestmann and B. C. Pierce. Decoding choice encodings. *Inf. Comput.*, 163(1):1–59, 2000.
- [17] L. Nielsen, N. Yoshida, and K. Honda. Multiparty Symmetric Sum Types. Full paper available at: http://www.thelas.dk/index.php/Symmetric_Sum_Types.
- [18] C. Palamidessi. Comparing the expressive power of the synchronous and asynchronous pi-calculi. *MSCS*, 13(5):685–719, 2003.
- [19] K. Prasad. Broadcast calculus interpreted in ccs upto bisimulation. In *Electronic Notes in Theoretical Computer Science*, volume 52, pages 83–100. Elsevier, 2001.
- [20] K. Takeuchi, K. Honda, and M. Kubo. An Interaction-based Language and its Typing System. In *PARLE'94*, volume 817 of *LNCS*, pages 398–413. Springer, 1994.
- [21] A. ten Teije, S. Miksch, and P. Lucas. *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*. Studies in Health Technology and Informatics. IOS Press, 2008.